Whitepaper | September 2010

radisys®

DPI: Deep Packet Inspection Motivations, Technology, and Approaches for Improving Broadband Service Provider ROI

CONTENTS

The Drivers

Wireless and fixed network broadband service providers face the challenge of increasing return on investment in the face of increasing infrastructure costs to keep up with unprecedented data demand by subscribers while opportunities for increasing revenues remain elusive. Service providers are faced with unprecedented demand for more and higher speed bandwidth in the face of new applications, increased video use, and subscriber growth. Network convergence where all services will be provided over IP on the same technology—whether as LTE in wireless or over a fiber/coaxial/copper connection in fixed—forces the question of "fair use" and prioritization. Voice, video on demand, live video, network DVR, multimedia messaging, videoconferencing, browsing, and later medical, home monitoring, energy use, smart grid, and other technologies both as a bundled offering with a service provider or their partners as well as over the top all share the same infrastructure in IP services to a subscriber. Considering IP is the glue for all communications whether for entertainment, education, employment and work, payment, and a new round of machine to machine applications to support health, convenience, home security, metering, and efficient energy use, service providers have a unique position to provide the environment—one that ensures good performance and security—that enables this convergence and innovation to all IP. Service providers need an awareness of the applications that use IP in order to provide the key roles of security and fairness and DPI is the cornerstone in enabling those through informed decisions. Key motivations for DPI can be categorized as follows:

- *Cost Reduction:* Through increased knowledge of capacity, performance bottlenecks, and as well through traffic management more subscribers can be served from the same infrastructure. Subscriber retention increases through better management of the network combined with proactive engineering to add capacity to address performance concerns.
- *Increased Revenues:* Through the offerings of tiered services to both subscribers and business partners. New services such as content filtering to limit access to content that is appropriate for a child or employee, network security services to detect and/or prevent malware, crimeware, and sensitive information loss, and services to business partners for increased ad revenues and/or Quality-of-Service guarantees to deliver traffic.
- *Emerging Regulatory Compliance:* Service providers are increasingly faced with lawfully required surveillance and/or assistance for law enforcement and/or national security/defense, requirements to offer content filtering, requirements to report access to unlawful content, and similar, such as interests in detecting access and transfer of pirated content.

The Definition

DPI expanded is "deep packet inspection" however what is often asked for is not DPI, but the capabilities DPI enables such as traffic shaping, admission, content access restrictions, information extraction about subscribers from their packet traffic, and so on and at a higher level the DPI applications which form the basis of useful services to a subscriber or capabilities to a service provider.

- DPI can be broadly defined as the ability to collect information and optionally take action based on the information in or that can be inferred from the content of the communication. The applications running over IP increasingly, 1) are Dynamic and Distributed. They change, are stateful and operate from multiple sources, 2) involve Protocols, require the ability to understand not just a packet but the communication and syntax across multiple packets and ports, and 3) may need Intervention, such as blocked if unauthorized, or shaped to provide reasonable throughput overall.

DPI generally implies broad capabilities as determined by the goals of the service provider. DPI technology is deployed in a service provider context, and as such, will often co-reside with other network functions that are already providing communications services to subscribers, such as access gateways, or are in the communications path of traffic, such as a router, border gateway, security or VPN gateway, firewall.

Radisys' Approach to DPI

Given the focus of DPI technology and applications by service providers, service providers have similar requirements as other network infrastructure residing in the network. Reliability, fault tolerance, safety, high capacity, efficient power utilization, etc. are the norm. DPI equipment can be expected to be placed in data, computing, and switching centers and as such must meet the requirements to be housed with other nearby equipment. Given advances in computing and packet processing, the ability to have independent but co-resident DPI applications with other packet processing intensive applications (such as security gateways, access gateways, and others) becomes feasible.

For many service providers and their suppliers, ATCA has emerged as the standard for open, reliable, with a large ecosystem of hardware and software suppliers that meets key service provider requirements. A variety of compute from traditional designs to massively parallel multicore, storage, I/O, plug-in mezzanine cards, dedicated content processing, string/pattern match by co-processors/accelerators, and packet processing options have emerged in ATCA that provide for numerous options depending on the DPI application requirements. A collection and balance of capabilities combined with high reliability, capacity, I/O, and wide variety of options that provide the key enabling technology are used to build leading DPI applications. Radisys' broad platform and leadership across compute and resource blades (CPU, NPU, DSP, etc.), 10G/40G I/O to processing blades, high-capacity switching and I/O capabilities, and a leading array of chassis options, all certified by a single supplier make Radisys ATCA the choice for DPI applications platforms. Several software partners supplying high-availability, system management, IP-layer foundation processing are available through Radisys to speed applications development.

Applications

DPI applications are quite diverse and span a broad range of technical segments and capabilities.

- *Policy enforcement* as traffic shaping and prioritization, access controls and admission, and content filtering attract the most attention and are at the heart of both fixed and wireless broadband network evolution given their obvious importance in providing fair use, the option for new services, and the ability to enhance subscriber experience through traffic management.
- *Network security* applications ranging from basic firewalls to network-based antivirus, intrusion detection and prevention, data leak prevention, anti-spam, and anti-SPIT & SPIM (spam internet telephony and spam instant messaging), web-application firewalls are emerging both in conjunction with endpoint security software (to address gaps where endpoint software simply cannot be updated fast enough to address an outbreak) and in network-based security approaches where devices may not have endpoint security capability (such as in embedded or machine to machine applications).
- *Network and subscriber analytics* applications are commonplace—these aid the service provider in gauging the overall health of the network by pinpointing performance and capacity as well as provide the service provider with a greater understanding of their subscriber's behavior that may be used to enhance marketing revenues.
- *Monitoring and interception*—whether for purposes of lawful intercept or other regulations or to support problem diagnosis in live systems—are major DPI technical segments.
- *Content optimization* through proxying and modifying content to better adapt content—such as by reducing still and video image quality, reformatting web pages for mobile devices, and other techniques—given bandwidth and device constraints have been applied to allow more users to enjoy content with acceptable performance than otherwise.

- *Billing and Metering* applications to count the volumes or rates of traffic, but with more complex schemes to account for a mixture of both free, partner, and paid traffic to support schemes where a subscriber's traffic may be capped to a certain volume, may be paying by the byte, or other schemes (such as payments between a service provider and content provider for certain types of traffic and/or application usage).
- *Content caching* applications where a service provider may elect to serve cached content via intercepting traffic and returning the content directly.
- *Application Distribution and Load Balancing.* These employ DPI technology to examine packet content (possibly far into the packet) and re-write the packet to direct the packet to a different destination for purposes of load balancing, fault tolerance, or other uses.
- *Modification and Injection* applications examine content and modify packet content for a variety of purposes, such as to insert tracking ids, rewrite packet headers (such as TOS), and may inject new packets or traffic as a result (for example: injecting TCP resets to interfere with P2P traffic).

Key Requirements

Examining the requirements of these key segments one can distill two key properties which fundamentally govern the overall design:

- Need to act in real-time
- Need to modify or intervene in packet content or processing

Many applications require real-time performance, with real-time being defined as the available time to make a decision as permitted by the subscriber or application. Several applications do not require real-time performance, such as subscriber and network analytics, intrusion detection security applications, email virus scanning, and some billing and metering systems depending on service provider desires; these applications are often not-time sensitive and data can be stored and processed at a later time. Other applications require real-time such as traffic prioritization and content filtering.

The need to modify or intervene and to be in the path of communications primarily influences the fault-tolerance and tolerable delay. Subscriber and network analytics, interception and monitoring, and some billing and metering applications can be viewed as passive. They do not necessarily need to directly modify communications. Policy enforcement, content optimization, caching, and others require modifying packet content, delaying, prioritizing or reordering packets, or blocking packets that clearly imply being in the communications path.

These two key properties of real-time and the requirement to modify the packet content govern key platform options and requirements. Applications which may not need to operate in real-time may choose to store and buffer and hence may have storage requirements or one may choose a storage-CPU tradeoff to extract key data and store only the metadata. Being within the communications path, leads to increased high-availability requirements. The DPI application must be robust against single points of failure and highly available to avoid contributing to overall communications loss. Many applications have both real-time and in-the-path requirements, such as policy enforcement. Placement in the path may tend to favor co-residency with existing network elements; however, operational and capacity considerations come into play, such as the extra capacity needed for the application and the upgrade and maintenance of the DPI application which may be frequent (such as needs to update threat detection capabilities quickly in response to new threats, or to act on new user applications and protocols as they become prevalent).

DPI in Mobile and Fixed Broadband Service Provider Networks

Given the nature of IP, DPI elements may be placed anywhere in the communications path, and may be co-resident with access gateways, routers, security gateways, border gateways and so on or may be separate elements. Special care must be taken for DPI applications that modify or intervene in packet content should the DPI element be lower in the communications stack than IP, such as in the radio access network or in a wireless packet core before a GGSN or P-GWY given these function beneath the IP layer and may also be subject to mobility and may conflict with billing, QoS, and lawful intercept design. DPI application functions can also be distributed and cooperate with inherent capabilities within the network, for example one element can be responsible for detecting events of interest (such as the use of P2P) and another element can be responsible for policy enforcement (such as blocking P2P or traffic shaping P2P). Increasingly GGSNs, CMTS, and LTE P-GWYs have inherent capabilities to block and/or shape traffic. Some DPI applications inherently require distribution, such as analytics applications that report the delay between points in the network. DPI elements often interact with policy control elements such as PCRFs for either policy decisions or to leverage a policy enforcement capability elsewhere in the network, such as on a GGSN or LTE P-GWY. Interactions with OSS or other administrative systems to notify of events or information for billing, trouble notifications, and violations for terms of service for using prohibited applications may also occur. DPI elements responsible for capturing and reporting events may also be part of other systems responsible for billing, subscriber and network analytics, interception, monitoring, or other systems. The following network diagram is representative of placement of DPI elements and functions and their relationships to common network elements and support systems.

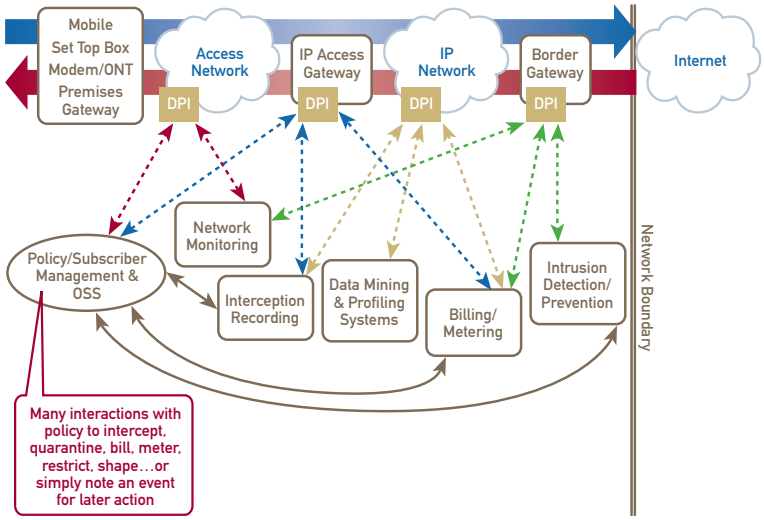

Figure 1. *Generic IP Network Transaction—User Terminal <-> Internet*

Key DPI Functions

DPI applications whether policy enforcement, security, or analytics have common underlying technical needs. Several key needs are identified below (Figure 1).

Reassembly

Packets may be fragmented, be reordered, retransmitted, duplicated, and so on. A fundamental capability is the ability to reassemble a communications session based on the packets observed as transmit/received over the wire. Enough of the session must be available to be able to characterize, extract key information, gather statistics, and so on.

Protocol Detection and Information Extraction

The ability to detect and discover a given protocol is in use during a communication is fundamental. Statistics about the protocols in use, permission to use a given protocol, billing or metering, traffic shaping, etc. are all reliant on knowing what protocols are active. Useful information is often present within a protocol, such as the subscriber's identity, filenames, identities or addresses of another party (telephone #, email address, SIP URI) that can be extracted.

This protocol detection capability is becoming more complex for some DPI applications considering the nature of many newer user applications use HTTP as a transport for others (such as SIP) both as a normal course as well as to evade detection…the ability to detect and extract information from a nested or embedded protocol must also be considered.

Pattern Matching

Pattern and/or string matching is often found in many DPI systems, such as to look for particular text or byte strings for a variety of reasons. Many protocols are text-based and can be readily identified by determining if certain keywords exist, similarly virus and malware are often identifiable by text or specific byte strings. Virtually all DPI applications incorporate some type of pattern matching technology, such as pattern matching by regular expressions. Regular expression based pattern matching has emerged in hardware both as co-processors and incorporated into packet processors.

L2-L4 Packet Header Processing

Although one may argue this is not "deep" inspection, the ability to rapidly characterize the packet headers is key. Subscriber identity can be inferred in some cases from an IP address, attributes can be attached to an IP address (common services, nature of services or content provided), and so on. Packet headers can serve as an input to classifying a communications, protocols often operate on well-known ports. DPI applications which modify content will often need to re-write TCP or other headers. Traffic management applications may re-write the TOS/DSCP header or 802.1p headers to ensure proper prioritization. Advances in hardware assisted packet processing may be leveraged by DPI applications.

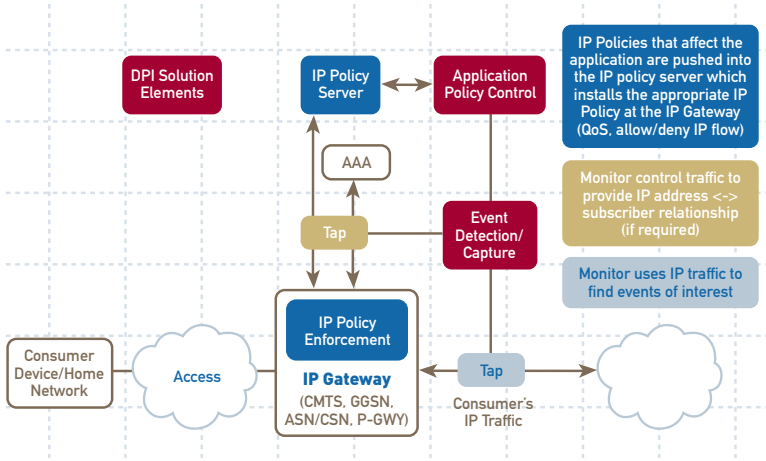

Figure 2. *Distributed Detection and Co-Operation with IP Policy Enforcement*

DPI Use Cases and Architectures

DPI applications can take many forms and there are a variety of placements and options to consider in any deployment. Is DPI integrated with another element that already provides packet processing? Can functions be distributed throughout the network? Can traffic shaping capabilities inherent in 3G and LTE wireless and newer cable systems be leveraged? What is the relationship to policy? Can network information be leveraged? What is the reasoning for integrating vs. external?

Approaches that integrate DPI into existing elements as well as external are illustrated below (Figure 2).

This first case is illustrative of DPI applications that augment policy enforcement capabilities that might exist in a newer deployed WiMax, PacketCable, or 3G network or LTE with a PCC or dynamic policy capability, as well as cases where the nature of the DPI application requires a complex, frequently

updated, or new capability based on monitoring but where an enforcement capability does exist in an existing node (e.g., such as a GGSN, CMTS, or P-GWY), such as traffic shaping and/or the ability to block traffic or where the ability to dynamically alter QoS for application traffic flows exists. Note that both subscriber traffic and network control plane traffic are monitored, the purpose of monitoring the network control plane is to be able to identify subscribers to IP address relationships to enable policies that can be directed toward individual subscribers. Both user and network traffic are observed to capture events, when an event of interest is detected (e.g., such as a prohibited application, or an application which might benefit from traffic management), the DPI application policy control is notified which then causes the appropriate IP policy to be applied at the IP gateway (Figure 3).

The second case illustrates packet or application policy enforcement being applied external of existing equipment, which may be the case of a retrofit of an existing network where dynamic policy capabilities do not exist and/or the nature of the DPI application does not well co-exist with existing gear (such as reformatting content, transcoding, injecting traffic. DPI applications which modify and inject traffic as well as applications where one might need to intervene to prevent traffic from passing, such as web content filtering or intrusion/network malware prevention systems). As in the first case, both subscriber and network traffic are monitored and appropriate events are notified to permit decisions about actions to take on a particular subscriber's traffic (Figure 4).

In the last case, consider a device containment system that is meant to contain faulty or infected devices or applications in a 3G network, and that both the detection and enforcement capabilities exist on a GGSN. The GGSN has awareness of both subscriber traffic and subscriber identity, no need exists to monitor network control traffic. DPI functions responsible for detecting traffic, applications, usage, and so on that may be consistent with a malfunctioning or malware infected device (or tethered PC) are co-resident on the GGSN,

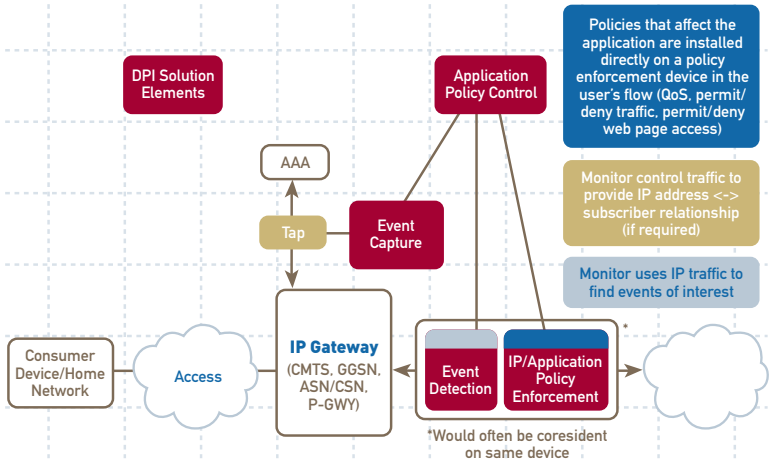

Figure 3. *Integrated Detection and Enforcement*

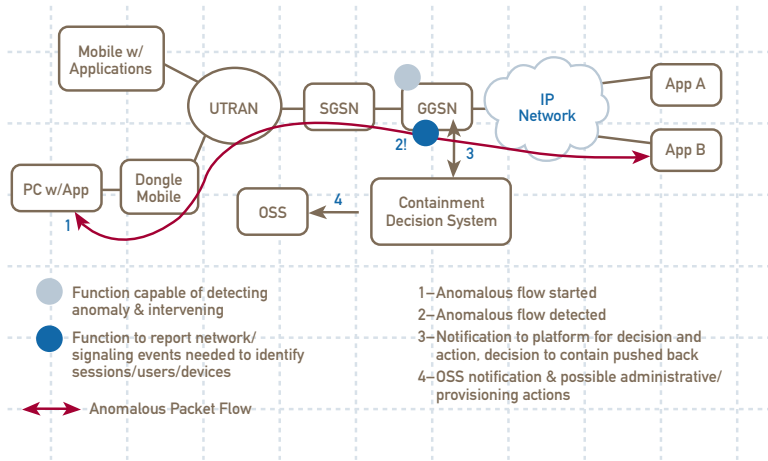

Figure 4. *Device Containment Example*

interfaces to report subscriber identity or other control events also exist. An external policy system—one which configures the events to monitor, the subscribers to monitor, and one that decides on the action to take should an event be reported—is used and may be proprietary given the state of 3GPP PCC standards (application specific policies and policy management while possible through application specific extensions to PCC are not yet defined in standards).

Radisys Solution for DPI Applications

Radisys provides a range of ATCA equipment from chassis, CPU, packet processing blades, switch and integrated system management blades, and access to a range of software and embedded partners to provide a solid foundation to build DPI applications. Radisys provides industry leading capacity and performance and is leading the charge to 40G platforms. As an example, a typical DPI platform configuration is shown in Figures 5 and 6.

More information on the ATCA CPU boards, NPU boards, AMC and Chassis can be obtained at: http://www.Radisys.com/Products/ATCA.html.

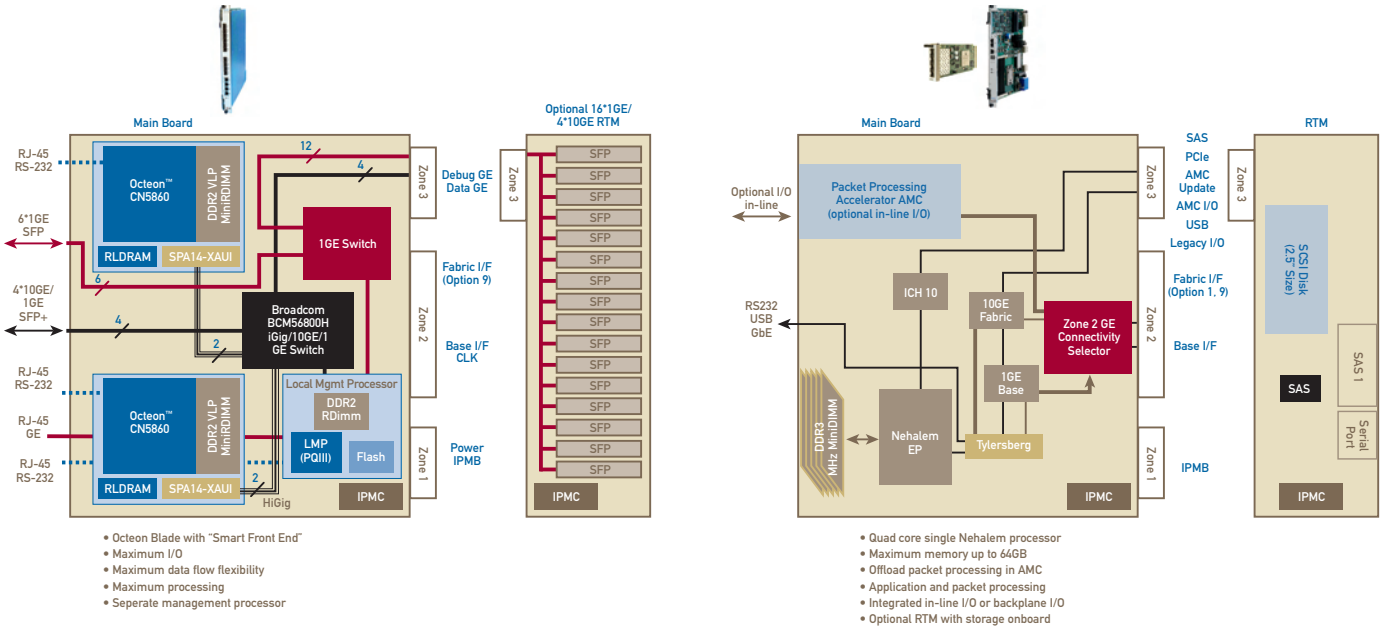

Figure 5. *ATCA Choices for DPI and I/O Resources*

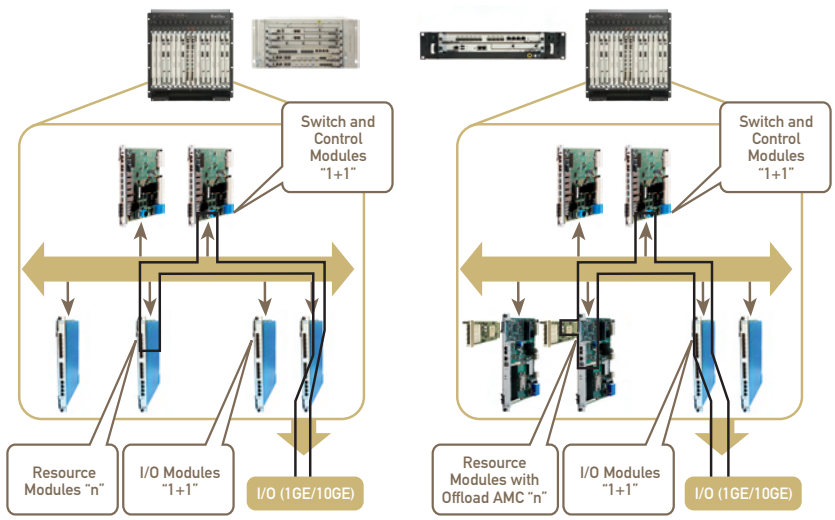

Figure 6. *Implementing DPI in ATCA and AMC Architectures*

Depending on the DPI application, an appropriate mix of CPUand/or packet processor blades may be used. DPI application providers may have substantial investment in existing code that is not easily ported to a packet processor or simply may not require absolute real time and packet forwarding capabilities needed for the DPI application, such as when the application is required to be in the path. The application may also be compute or storage/memory intensive. Packet processing capabilities in CPU-based systems are also improving and benefit from overall advances in CPU design. Radisys stays ahead of the curve by offering the latest generation of Intel® processors and provide expandability options by leaving room for an AMC which can be used for storage or a co-processing application.

Dedicated packet processors have clear use in DPI applications, especially when the application is in-line and latency is a consideration given both packet forwarding and DPI requirements must be satisfied. A growing number of packet processors integrate a variety of hardware acceleration capabilities, such as on-board regular expression pattern matching capabilities which can be used for protocol detection and threat detection; decompression/compression capabilities useful to examine content given much content (files/images/script/xml) may be compressed; and encryption/decryption capabilities (which may be useful in examining content where a service provider may have session keys) which can be leveraged to improve both speed and throughput. A key consideration for any in-line DPI application is the amount of time allowable for examining, classifying, and taking a decision on a packet. One cannot unduly delay a packet or cause significant jitter considering the increasing use of conversational voice and video over IP.

One should also consider the possibility of hybrid CPU-based and packet processor based systems. One may wish to leverage packet processor wire-speed, L2-L4 header processing and matching, and other hardware assisted acceleration or pattern matching capabilities to make fast decisions while allowing for other decisions or further analysis of communications by CPU blades where more complex logic may be executed and/or taken off path. CPU blades may also be employed in a packet processing system simply to allow for more division of work, with the packet processing blades being used for packet forwarding and packet inspection, and CPU blades are used for local policy decisions, performing off board queries, handle configuration or subscriber databases, reporting events, capturing and logging events, and so on.

A variety of 3rd party AMC co-processors that may be employed on CPU blades or on AMC carrier blades are available in the marketplace, including a Radisys AMC employing an Octeon® Plus processor and an emerging breed of what are termed “content processors.” Content processors seek to enhance certain classes of DPI applications by providing very large regular expression and pattern matching capabilities such may be useful for content classification applications, applications which need to execute large numbers of string searches in parallel, such as network antivirus prevention or detection systems, or content filtering applications, or have significant numbers of complex regular expressions to evaluate. Content processing perhaps provide an avenue for CPU-based applications to gain an option for performance improvement while maintaining an overall portable architecture.

Complete DPI applications must be assembled into an overall solution including OAM, system management, fault tolerance, packet routing, protocol stacks, database, and other capabilities, upgrade, logging, and DPI capabilities to detect and extract information from packet flows, provide utilities to perform pattern matches, shape traffic, and gather a variety of statistics, such as traffic volume, frequency of application usage and response time. Many software and system elements are integrated to form a DPI application. Radisys has several partners to provide system management, high-availability and fault tolerance frameworks, protocol stacks to support co-resident applications as well as lower level capabilities such as packet forwarding and traffic shaping. DPI toolkits to leverage the Octeon® architecture are available from Cavium. Radisys continually evaluates its software partner portfolio to add new partners & capabilities as demands evolve.

Summary

"DPI" is already being implemented to provide a range of services from content filtering, to web application firewalls, application aware policy enforcement, and others. DPI applications solutions are a mix of both compute and packet processing, real-time and non-real time processing, as well as in-the-path/on-the-wire and off-the-wire/off-the-path processing depending on the DPI application. As advances in computing and packet processing occur, the ability to inspect and act on more traffic to enable a variety of useful services—content-based traffic prioritization, network hosted intrusion and malware prevention systems, dynamic content filtering, and web application firewalls to address increasing mixtures of threats in "web 2.0" spanning image, text, script, java, flash, and other executeable objects—is enabled within the network. Faster availability of subscriber analytics enables better planning as well as can enhance a service provider's ability to participate in advertising revenues. Network analytics enable the service provider to understand the performance of the network and react sooner and more effectively to true capacity and equipment limitations.

With improvements in compute and packet processing performance in the same power/form factor occurs, more applications can be made co-resident enabling service providers to offer more services on the same infrastructure with better cost, reliability, and performance than through deployment of external solutions. External solutions though do have a role, and the combining or co-residency of a DPI application with a core function must consider other factors, such as the need for frequent updates required by some DPI applications to stay current whereas the core function may be very stable and require only routine patching and upgrading, some DPI applications may require updates several times a day whereas core network functions may go months between upgrades or patches.

In many respects, aspects of DPI are already ingrained in both wireless and newer fixed networks and will absolutely be required going forward given the convergence to all IP. Voice and critical applications require dynamic prioritization under all IP, and this capability is being built into LTE, WiMaX and next generation fixed networks infrastructure. This capability can be leveraged to enhance subscriber experience by adding applications knowledge and intelligence to the decision to prioritize and ensure good performance and as well the ability to *deprioritize* can be used to better manage limited resources.

As DPI applications are increasingly deployed in the service provider networks, there is a default expectation these applications will have the same or better availability, reliability, and quality as other parts of the service provider's offerings. Certain DPI applications, such as policy enforcement, are becoming integral to the overall service. The platform used to deploy service provider DPI applications must be carrier-grade and reliable, leading many to choose ATCA given its success in the service provider arena. Leveraging ATCA assures a wide variety of suppliers, long-life, and community of support.

Radisys offers a sound foundation for constructing ATCA DPI-based applications through a variety of compute modules and AMC options and packet processing blades as well as pre-certified HA, systems management, load balancing, routing, switching, QoS management, and protocol software toolkits to enable the fast and cost-effective creation of DPI ready applications.

radisys®

Corporate Headquarters
5445 NE Dawson Creek Drive
Hillsboro, OR 97124 USA
503-615-1100 | Fax 503-615-1121
Toll-Free: 800-950-0044
www.radisys.com | info@radisys.com

intel Embedded Alliance Premier